

## Re:fcntl(F\_GETLEASE) semantics??

**Source:** <http://linux.derkeiler.com/Mailing-Lists/Kernel/2005-08/3117.html>

---

**From:** Trond Myklebust ([trond.myklebust\\_at\\_fys.uio.no](mailto:trond.myklebust_at_fys.uio.no))

**Date:** 08/11/05

To: Michael Kerrisk <[mtk-lkml@gmx.net](mailto:mtk-lkml@gmx.net)>  
Date: Thu, 11 Aug 2005 08:49:12 -0400

to den 11.08.2005 Klokka 14:27 (+0200) skreiv Michael Kerrisk:

- > *And I pointed out that the existing behaviour (which is*
- > *still current in 2.6.13-rc4) is inconsistent:*
- >
- > <http://marc.theaimsgroup.com/?l=linux-kernel&m=111511455406623&w=2>
- >
- > *Some further testing showed the following (both open()*
- > *and fcntl(F\_SETLEASE) from same process):*
- >
- > *open() | lease requested*
- > *flag | F\_RDLCK | F\_WRLCK*
- > -----+-----+-----
- > *O\_RDONLY | okay | okay*
- > *O\_WRONLY | EAGAIN | okay*
- > *O\_RDWR | EAGAIN | okay*
- >
- > *In other words, a process can open a file read-write, and*
- > *can't place a read lease, but can place a write lease!*
- > *That does not seem to make any sense to me.*

Then what do you think that leases are supposed to do, and why?

AFAIK, the whole point here is to provide a method to allow CIFS and NFSv4 clients to be notified if there is some behaviour on the server that screws with the ability to cache data.

An exclusive (i.e. write) lease should mean that *nothing* other than your process is accessing the file. A client may cache the file data, metadata and read/write locks because nobody else can change that information, and nobody else holds locks on the file. It may also cache file acl/access information, and hence cache new OPEN calls.

A shared (i.e. read) lease means that there are currently no processes that can change the data or metadata (including your own). A client may cache data, metadata and read locks since there are no writers, and there is nobody holding write locks. The client may again cache OPEN calls as long as they are read-only.

Linux-Kernel: Re: fcntl(F\_GETLEASE) semantics??

Note that the kernel is still incomplete w.r.t. notification of changes. Holders of the oplocks/delegations need to be notified if the file is renamed, or if the acl/access information changes, say.

Cheers,  
Trond

—

To unsubscribe from this list: send the line "unsubscribe linux-kernel" in the body of a message to majordomo@vger.kernel.org

More majordomo info at <http://vger.kernel.org/majordomo-info.html>

Please read the FAQ at <http://www.tux.org/lkml/>